arXiv:1612.00565v1 [cs.RO] 2 Dec 2016

Programming by Demonstration with User-Specified Perceptual Landmarks

Justin Huang and Maya Cakmak

***Abstract*— Programming by demonstration (PbD) is an effective technique for developing complex robot manipulation tasks, such as opening bottles or using human tools. In order for such tasks to generalize to new scenes, the robot needs to be able to perceive objects, object parts, or other task-relevant parts of the scene. Previous work has relied on rigid, task-specific perception systems for this purpose. This paper presents a flexible and open-ended perception system that lets users specify perceptual "landmarks" during the demonstration, by capturing parts of the point cloud from the demonstration scene. We present a method for localizing landmarks in new scenes and experimentally evaluate this method in a variety of settings. Then, we provide examples where user-specified landmarks are used together with PbD on a PR2 robot to perform several complex manipulation tasks. Finally, we present findings from a user evaluation of our landmark specification interface, demonstrating its feasibility as an end-user tool.**

I. Introduction

Programming by Demonstration (PbD) is a user-friendly technique for programming complex manipulation actions on a robot, simply by walking it though the steps of the action. Previous work has demonstrated that a wide range of object manipulation actions can be represented as a sequence of end-effector poses relative to detected objects or other points of interest (*i.e.*, "landmarks") in the scene [1], [2], [3]. However, the expressivity and generalizability of this action representation strongly depends on the robot's ability to robustly perceive landmarks, as well as the granularity of those landmarks.

Unfortunately, previous PbD systems include rigid perceptual systems that limit and pre-specify the types of landmarks that can be referenced in the task. For instance, one system uses objects detected on a tabletop as landmarks for various manipulation actions [1]. This system will not work for objects on shelves, in clutter, or in scenes that have no horizontal surfaces, such as a doorknob on a door. In addition, it cannot localize specific *parts* of an object. A different system involves actions relative to object parts as landmarks, *i.e.*, ends of a rope to be tied [3], but these parts are detected with a special-purpose perception system that cannot be used for any other object. Another system involves easy-to-detect fiducials (also known as Augmented Reality tags) attached to object parts that are to be assembled by the robot [4]. This requires attaching fiducials to every single object or scene element that the robot might interact with.

In this paper, we address these limitations by allowing task-relevant landmarks to be specified during the demonstration. We refer to these as *custom landmarks*. We present a simple interface for specifying custom landmarks and an algorithm for localizing them in a new scene. We first characterize the algorithm's ability to localize landmarks in different scenes. Then, we demonstrate how manipulation actions programmed on a real robot using custom landmarks can generalize to different scenes. Finally, through a small-scale user study, we demonstrate the feasibility of our tool for use by novices. By contributing a flexible and open-ended perceptual system, our work allows manipulation actions that deal with a wide range of objects, object parts, or scene elements in very different environments to be programmed by end-users.

The authors are with Computer Science and Engineering Department, University of Washington, Seattle, WA 98195, USA.

II. Related work

This paper presents a perceptual system that expands the set of actions that can be programmed by demonstration. As such, it builds upon previous work in programming by demonstration, as well as in object localization and point cloud registration.

A. Programming by Demonstration

PbD is a widely used technique for programming industrial robots, with production-level systems that ship with robots like Kuka's LWR, ABB's Yumi, or ReThink's Baxter, among others. It has been an active research area since the 80s, resulting in several surveys on the topic [5], [6], [7]. While early research explored various representations for encoding and reproducing robot manipulator motions [8], [9], [10], [11], recent work has addressed challenges in manipulating objects [3], [1], [12], learning high-level task structures [13], [14], [2], and learning from non-expert demonstrators [11], [15]. PbD is also closely related to Learning from Demonstration, in which the robot's actions are typically represented as policies in a Reinforcement Learning framework [6].

Previous research on PbD has not focused on scene understanding, instead avoiding the perceptual challenges by using fiducials [4], marker-based motion capture [16], or simulated environments [2]. Others use special-purpose or limited perceptual systems such as object detectors on a flat, uncluttered tabletop surface [1], [17]. In contrast, our work focuses on extending the perceptual component of PbD. In one work, Ehrennmann et al. [18] compared object recognition techniques for PbD. These techniques were 2D image processing systems that recognized objects with known CAD models available. In our work, we localize 3D landmarks without assuming that CAD or other models of the landmarks exist.

B. Object localization

Our paper contributes an algorithm for localizing custom landmarks in new scenes, which builds on the iterative closest point (ICP) algorithm [19]. We represent custom landmarks using a point cloud and attempt to register the point cloud to part of the scene. Researchers have presented related algorithms to localize objects in a scene. Rusu et al. introduced an algorithm called SAC-IA [20], which computes an alignment between two point clouds by repeatedly matching 3D features sampled from the point clouds. Buch et al. [21] present a similar method with different features and geometric checks to optimize speed. One way in which our approach differs is that our landmark representation also defines areas of empty space which are expected to be unoccupied around the point cloud. This allows users to make distinctive landmarks out of otherwise nondistinctive point cloud segments (Fig. 1(a), 1(b)).

The problem of localizing landmarks can also be viewed as an object detection and recognition problem, which has seen major progress in recent computer vision research [22]. However, these systems are hard to apply to a user-friendly PbD system because they recognize a predefined set of objects and require large amounts of training data to work well. Additional work has gone into building representations of 3D shapes using deep neural networks [23], [24]. These representations could be used to compare sampled volumes from the scene to a custom landmark. However, the current resolution of these representations, $24\times24\times24$ or $32\times32\times32$ voxel grids, may be too limited for use in PbD.

III. Custom landmarks

A. Programming by Demonstration with landmarks

Our system builds on the PbD system by Alexandrova et al., implemented on a dual-arm PR2 robot [1]. Actions are represented as a sequence of 6D end-effector poses relative to the robot's base or to a *landmark*, *i.e.,* a point of interest in the environment that can be perceptually detected by the robot. Each pose also encodes the gripper's state (open or closed). Actions are programmed by first asking the robot to detect landmarks in the environments, then kinesthetically moving the robot's arms to desired poses, changing the gripper states if needed, and saving the pose using a verbal command. A pose becomes relative to a landmark if it is within a certain distance to it; otherwise, it will be relative to the robot's base. The user can later change the landmark associated with each end-effector pose through a graphical interface.

After programming, the robot executes an action by first localizing all the landmarks in the scene, which may have moved since the demonstration. The system in [1] accomplishes this by performing tabletop segmentation and returning segmented objects as landmarks. The robot then recomputes end-effector poses defined relative to these landmarks. Finally, it moves through the poses with both arms, opening or closing its grippers as needed. An action cannot be executed if a landmark is missing or if a pose is out of reach.

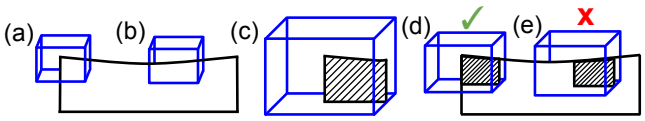

Fig. 1: An example that illustrates custom landmark specification and search. (a) A landmark representing the top left corner of a chair is selected by drawing a box around the corner. The top and left halves of the box are empty space. (b) A box with empty space above the chair, but not on the sides, represents any part of the top of the chair. (c) A close-up view of the landmark captured in (a). The shaded region represents the captured point cloud. (d) The corner landmark in (c) will match well with the corner of a new chair. (e) However, it will not match well when aligned with the top middle of the chair. Although all the points in the landmark (shaded region) match well with the chair, some of the chair is inside the rest of the box, which is expected to be unoccupied.

B. Custom landmark representation

In our system, landmarks are represented by a point cloud that captures the shape of the landmark, as well as a box that surrounds the point cloud. The boundaries of the box specify margins of empty space around the landmark that are expected to be unoccupied. The use of empty space is illustrated in Fig. 1.

The user interface for creating landmarks shows a point cloud view from the robot's depth sensor, as well as a 3D box-shaped selector, with controls to set its position and dimensions (see Fig. 3). The user moves and resizes the box to surround the point of interest, potentially including margins of empty space on the side. When the landmark is saved, the system records the subset of the scene within the box, as well as the pose and dimensions of the box. The interface provides a box-shaped selector aligned with the robot base for simplicity. However, a more advanced interface could be used to define an arbitrary shape and orientation for the custom landmark. The landmark is captured from a single view and from a single point cloud, and does not include color information. Future implementations could fuse multiple point clouds across time or viewpoints.

C. Landmark search algorithm

To make use of custom landmarks, we need a way of localizing them in a new scene. To describe our localization algorithm, we first formally describe the inputs and outputs. A point p is a location vector $p = (p_x, p_y, p_z)$ and a scene S is simply a set of points. A landmark ℓ is a tuple (P, B), where P is a set of points (*i.e.*, points in the point cloud selected by the user) and B represents a box, specified by a 6-dimensional pose and a 3-dimensional size vector.

Our search algorithm takes as input a scene S, represented by the complete point cloud captured before an execution, and a landmark ℓ to be localized in that scene. It outputs a set of landmarks O, where each landmark $\ell_o \in O$ is potentially an instance of the input landmark ℓ in the scene S.

Algorithm 1: FindLandmark

Input : *scene*, *landmark*, miscellaneous *parameters*
Output: a list of aligned landmarks in the scene, or empty list if not found

```
1  crop scene;
2  downsample scene using a voxel grid;
3  samples = randomly sample points from scene;
4  candidates = [];
5  for (sample in samples) {
6      c = copy of landmark;
7      move c such that c.cloud is centered on sample;
8      c = run ICP to align c.cloud with scene;
9      c.error = CandidateError (c, scene);
10     add c to candidates;
11 }
12 remove all candidates that do not have the lowest error
     within a certain radius (non-max suppression);
13 output = {c ∈ candidates | c.error < threshold};
14 return output;
```

Algorithm 2: CandidateError

Input : *scene*, *candidate* landmark
Output: error score

```
1  error = 0;
2  denominator = 0;
3  croppedScene = scene cropped to candidate.box;
4  visited = [];
5  for (scenePt in croppedScene) {
6      candidatePt = nearest point in candidate to scenePt;
7      error += distance between scenePt and candidatePt;
8      denominator += 1;
9      add candidatePt to visited;
10 }
11 for (candidatePt in candidate.cloud) {
12     if (candidatePt not in visited) {
13         scenePt = nearest point in scene to candidatePt;
14         error += distance between scenePt and
             candidatePt;
15         denominator += 1;
16     }
17 }
18 return error / denominator;
```

Pseudocode for the algorithm is given in Algorithm 1. First, we crop and downsample the scene according to application parameters.[1] Then, we randomly sample scene points and initialize an instance of ℓ at each sampled point. Next, we run the ICP algorithm to align the landmark's point cloud P with S, which produces a modified landmark ℓ'. For each landmark ℓ', we compute an error metric (Algorithm 2) using ℓ' and S. We then perform non-max suppression so that we do not produce duplicates of the same result. Finally, we filter results by thresholding on the error metric.

The error metric (Algorithm 2) can be thought of as the mean distance between the points of the scene (within the landmark box), and the nearest points of the candidate landmark (lines 5-10), and vice versa (lines 11-17). Adding a margin of empty space around the landmark helps eliminate false positive matches. If scene points are found where there is expected to be empty space, then they will increase the mean error, since the nearest points on the landmark are far away. This process is illustrated in Fig. 1(d) and 1(e).

Having all the points of the scene (within the landmark box) match well with the landmark does not imply that all the points of the landmark match well with the scene. This comes up when the landmark is aligned with a small part of the scene, such that there are more landmark points than scene points. To avoid false positives in this scenario, we ensure that each point of the landmark contributes to the error score at least once (Algorithm 2 lines 11-17).

D. Advantages and limitations

The perceptual system presented in this paper is flexible enough to represent objects, parts of objects, or parts of the scene. It has the following advantages, compared to existing perceptual systems for PbD:

1) Does not assume a tabletop scene: Users demarcate landmarks wherever they are in the scene. A custom landmark can represent an object resting on a table, sitting on a shelf, mounted on a wall, etc.

2) Can localize landmarks in hard to segment scenes: When objects are placed too close together, object segmentation algorithms can easily confuse them for a single object. The landmark search algorithm, however, searches for custom landmarks uniformly throughout the scene. This makes it possible to localize a landmark in contact with another object, as long as it remains visible.

3) Can represent partly occluded objects: Even systems that use full object models for landmarks can have a hard time localizing partly occluded objects. For example, when bowls are stacked, only the rim of the top bowl may be visible, making it hard to match to a model of the bowl. A custom landmark can be used to represent just the bowl rim, allowing the robot to localize the top bowl of the stack.

4) Can generalize across objects with the same part: Custom landmarks can also be used to represent a part of an object. For example, while different human tools with similar handles (*e.g.*, feather duster, sweeper, squeegee) would normally require a separate object detector for each tool, a custom landmark that represents only the handle could be used to detect all of them and easily transfer actions demonstrated for one tool to the the other.

5) Does not require task-specific detectors: Our system localizes custom landmarks based on their shape and is agnostic to the task at hand. This allows it to be used for a variety of tasks that might otherwise require custom

[1] For this paper, the scene was cropped to a volume in front of the robot roughly equivalent to the reach of its arms. The scene was downsampled to a leaf size of 0.005 meters. We sampled 5% of points in the scene, up to a maximum of 1000 points. Our non-max suppression radius was 0.03 meters. The threshold for our error metric was set to 0.0055 meters.

TABLE I: Scenes and landmarks in the *simple* setting.

Scene	Landmarks
A bowl and a cup on a table	Bowl, cup
A Tide bottle on a table	Tide bottle
A Tide bottle and a spray bottle	Tide bottle, spray bottle
Three balls on a table	Balls

perception systems (*e.g.*, handle detectors, bowl or bowl rim detectors). This is also useful for fixed parts of the scene. For example, a custom landmark can be used to localize some prominent feature of a laundry control panel, such as the central dial, or a corner of the machine. The robot can then press different buttons on the panel, whose positions are known as offsets from the landmark.

6) Does not assume "objectness": Custom landmarks can be arbitrary parts of a scene, such as a drawer handle, the top slot of a recycling bin, or a window sill. They can even represent the lack of objects, *e.g.,* using a flat horizontal patch to search for an empty spot on a cluttered surface.

The main limitation of our system is that it relies on landmarks having a unique shape in the scene. In semi-structured environments, it could be the case that the scene is designed to not distract from the landmarks, but this is not true in general. Custom landmarks also do not make use of color information, so they can't distinguish between different colors of the same object, or recognize color patterns. Because we only capture custom landmarks from a single view of a point cloud, it can be brittle to viewpoint changes, such as when a non-symmetric object is rotated. Finally, our system puts the burden on users to intelligently create the custom landmarks to maximize their uniqueness. To provide good landmarks, users must understand how Algorithm 1 and 2 work, which requires explanation and training.

IV. Search algorithm experiments

Our first evaluation characterizes the search algorithm's ability to localize landmarks in a variety of settings.

A. Settings

The algorithm was evaluated in three kinds of settings: *simple*, *cluttered*, and *unconventional*. Different sets of landmarks were created for each setting (Fig. 2).

In the *simple* setting (Fig. 2(a), 2(b)), objects were spaced apart on a tabletop and the landmarks represented whole objects. These scenes, detailed in Table I, were representative of scenes used in previous PbD work and showed how our system could function in lieu of tabletop segmentation.

In the *cluttered* setting, detailed in Table II and illustrated in Fig. 2(c) and 2(d), objects were in contact or occluding one another. Additionally, one of the scenes was a shelf scene, rather than a tabletop. Some of the landmarks represented parts of objects rather than the whole object itself. These scenes were designed such that tabletop segmentation would not have worked.

The *unconventional* setting (Fig. 2(e), 2(f), Table III) did not contain any tabletop scenes, and landmarks represented parts of the scene (*e.g.*, a drawer handle) or unique objects like Griples.[2] This was designed to demonstrate how custom landmarks can handle recognition tasks that might otherwise require a custom, task-specific detectors.

TABLE II: Scenes / landmarks in the *cluttered* setting.

Scene	Landmarks
A stack of bowls and a cup on a table	Rim of top bowl, cup
Two cups and stacked bowls placed on two different levels of a shelf	Rims of top bowls, cups
Bowls placed in contact, containing items including a tennis ball, on a table	Rims of bowls, tennis ball
A Tide and a spray bottle in contact on a table	Tide bottle, top of spray bottle
A spray bottle, occluded on bottom	Top of spray bottle
A different spray bottle with tennis balls and other clutter, on a table	Top of spray bottle, tennis balls

TABLE III: *Unconventional* scenes and landmarks.

Scene	Landmarks
Chair in front of robot	Corners of chair backrest
Lowered chair in front of robot	Corners of chair backrest
Partly full Griple tool rack	Griples, Griple mounts
Full Griple tool rack	Griples, Griple mounts
Empty Griple tool rack	Griples, Griple mounts
An under-desk drawer	Drawer handles
The under-desk drawer with bottom drawer open	Drawer handles

B. Measures

The localization algorithm was run for each scene and landmark pair per setting. The correctness of the output was hand-labelled by the first author. We report the precision and recall of the algorithm under three different assumptions:

1) All landmarks are searched for in all scenes.
2) Only certain landmarks are searched for in certain scenes (*i.e.*, they are *in context*), as enumerated in Tables I, II, and III. This can be a valid assumption in real-world settings. For example, if the robot goes to a Griple tool rack, it is probably either searching for a Griple or a Griple mount, but not a chair.
3) In addition to assumption 2, if multiple instances of a landmark are in a scene, we only want to localize one of them. This is a valid assumption for PbD actions, which only operate on one landmark at a time.

C. Results

Precision and recall values for the evaluation settings described above are summarized in Table IV. The results

[2] Griples are 3D-printed adaptors that can be attached to tool handles to make them easier for the PR2 to grasp, see [25]. Tools with Griples can be hung from 3D-printed Griple mounts (Fig. 2(f)) fastened to a wall.

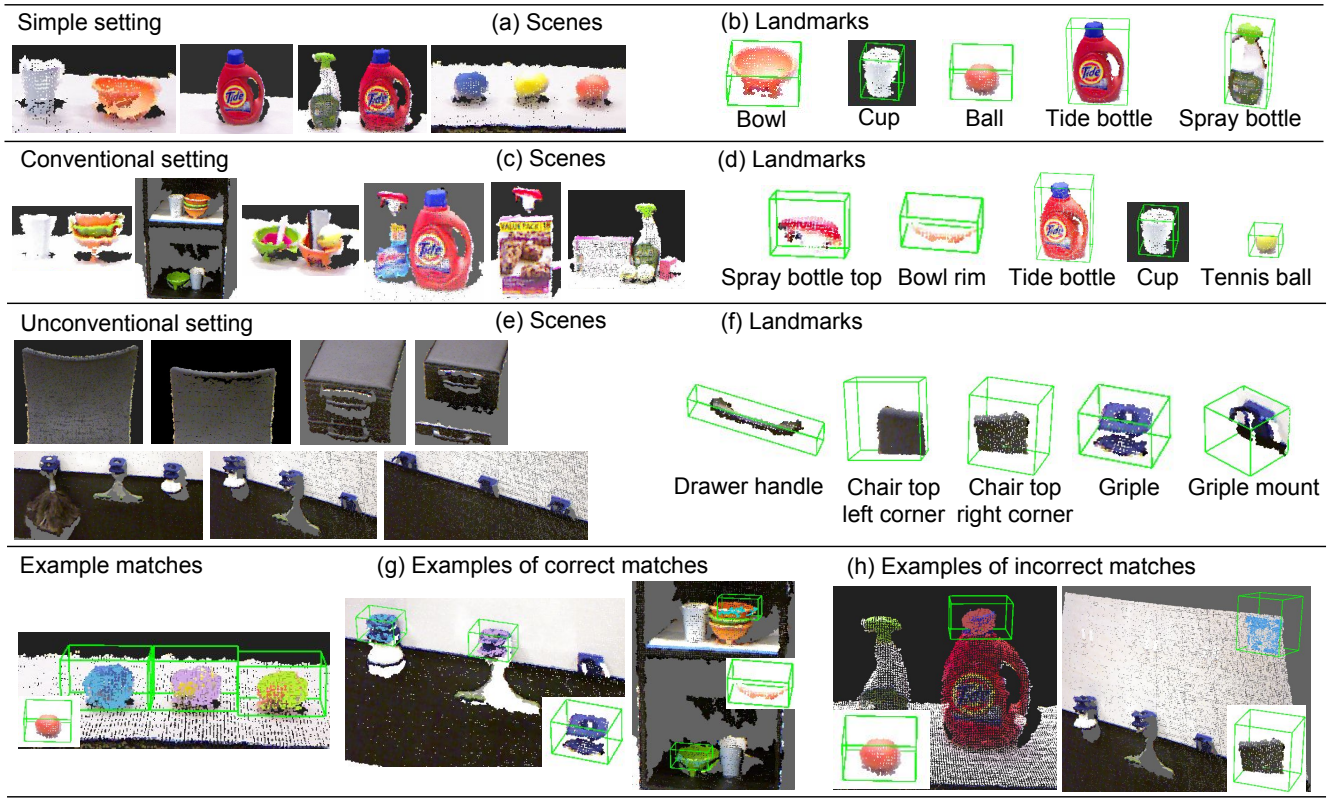

Fig. 2: (a) Scenes in the *simple* setting were such that objects could be easily segmented. (b) Landmarks in the *simple* setting all represented whole objects. (c) Scenes in the *cluttered* setting had objects placed in contact or stacked on each other. This setting also included a shelf scene. (d) Some landmarks in the *cluttered* setting represented parts of objects, such as the top of a spray bottle or the rim of a bowl. (e) Scenes in the *unconventional* setting were not on tabletops. (f) Landmarks in the *unconventional* setting were unique objects like Griples or Griple mounts. Or, they were part of the scene, such as a chair corner or drawer handle. (g) Our system worked in simple scenes. It also found the two bowl rims on the top and bottom shelves of the shelf scene, and found unique landmarks like Griples. (h) Our system incorrectly confused the cap of a Tide bottle with a ball, and found a chair corner at the corner of a vertical surface.

show that our system works better in situations where landmarks are searched for *in context*. The recall rate is further increased when the system only needs to find one instance of a landmark—this shows that the system will often find at least one instance of a landmark, even though it may not find them all. Our system is least successful when searching for arbitrary objects in arbitrary scenes, which results in a larger number of false positives.

In the *simple* setting, our system mistakenly localized a ball at the cap of a Tide bottle. This is shown in Fig. 2(h). This error occurred in both scenes which had the Tide bottle in it. The system localized all other objects in all other scenes in this setting correctly.

In the *cluttered* setting, our system had the same issue with a tennis ball being localized at the cap of a Tide bottle. It also mislocalized the spray bottle top to the back of the bottom shelf in the shelf scene. The back of the shelf had an abnormal appearance due to the robot not having full depth data. This shows that, in our system, the smaller and less constrained the landmark, the more likely it is to be matched incorrectly in the scene. However, the system was able to successfully localize objects in cluttered and partly occluded configurations. For example, it was able to localize a bowl rim in a stack of bowls, including in the shelf scene (Fig. 2(g)). In the scene where a spray bottle was in contact

TABLE IV: Precision / recall scores of the landmark search algorithm, under the assumptions from Section IV-B.

***Simple* setting assumption**	**Precision**	**Recall**
All landmarks searched in all scenes	72.72%	100%
Landmarks and scenes in context	100%	100%
Only find one instance of landmark	100%	100%
***Cluttered* setting assumption**	**Precision**	**Recall**
All landmarks searched in all scenes	75%	60%
Landmarks and scenes in context	100%	60%
Only find one instance of landmark	100%	81.82%
***Unconventional* setting assumption**	**Precision**	**Recall**
All landmarks searched in all scenes	64.71%	64.71%
Landmarks and scenes in context	100%	64.71%
Only find one instance of landmark	100%	80%

TABLE V: Demonstration Task 1: Hanging tools

Scene and objects: Three Griple mounts on a vertical surface and three tools with Griples that could hang from the mounts **Landmarks:** Griple, Griple mount **Actions:** Picking up tools from the rack and hanging them back	
Test scenarios	***Success?***
Picking the only tool from a rack (Fig. 3(a))	*Yes*
Picking a tool from a full rack	*Yes*
Picking a tool with no tools available (should not act)	*Yes*
Placing a tool on an empty rack	*Partial*
Placing a tool on an full rack (should not act)	*Yes*

TABLE VI: Demonstration Task 2: Grasping a chair

Scene and objects: The back of an office chair facing the robot **Landmarks:** Top of the chair backrest **Actions:** Grasp the backrest of the chair with two arms	
Test scenarios	***Success?***
Five different scenes with the chair raised to different heights, placed in different positions and orientations	*4/5*

TABLE VII: Demo Task 3: Picking bowls from a shelf

Scene and objects: Similar to the shelf scene in the *cluttered* setting (Section IV-A), with bowls in contact with a paper cup. Scene-1: one bowl on the top and one bowl on the bottom shelf; Scene-2: two bowls stacked on the top shelf; Scene-3: only cups. **Landmarks:** Bowl rim **Actions:** Picking up a bowl from its rim	
Test scenarios	***Success?***
Picking the two bowls in scene-1	*2/2*
Picking the two bowls in scene-2 (Fig. 3(c))	*2/2*
Picking up a bowl in scene-3 (should not act)	*Yes*

with a Tide bottle (Fig. 2(c)), the system was able to localize both bottles individually, whereas a tabletop segmentation algorithm would consider them to be one large object.

In terms of recall, the system failed to find some expected landmarks. In one scene, we put a tennis ball inside a bowl, to see if the system would find either the tennis ball or the bowl rim. Both objects intruded on the expected empty space regions of the other, so the system found neither.

In the *unconventional* setting, the only false positive was to localize the corners of a chair at the corners of a whiteboard, from which we hung Griples. This error was repeated and counted multiple times in our data, for the left and right chair corners, and for all three scenes using the Griple tools. It also failed to find the drawer handles in one of the two drawer scenes. However, the system did manage to localize objects like Griples, Griple mounts, and chair corners in the appropriate scenes.

D. Discussion

While our results demonstrated that our perception system can perform what it was designed to do as part of a PbD system, it was not robust in the general perception problem of finding arbitrary objects in arbitrary scenes. In particular, small objects are more likely match unexpected parts of the scene (*e.g.*, the ball being matched to the cap of the Tide bottle). Distinctive landmarks, such as those with empty space in its interior, work better. For example, although the bowl rim landmark is small, our error metric will not confuse it with most areas of the scene, because it expects to find empty space inside bowl rim's proposed location. Inversely, the chair backseat corners, which had little empty space inside, were confused with the corners of the whiteboard. Although the intent of the chair corner landmark was to only match with other chairs, it will in fact match with the corner of any flat surface. This suggests that users should try to design custom landmarks to be unique-looking. For example, if it can be assumed that balls are on a tabletop, then it could help to include part of the tabletop beneath the ball as part of the landmark. This would avoid the ball being confused with the similarly-shaped Tide bottle cap.

V. Demonstrations on a PR2 robot

This section provides several examples of how custom landmarks can be put into practice within PbD. We describe six complex manipulation actions that we programmed by demonstration using custom landmarks. These actions were not possible using previous systems such as in [1]. An important aspect of these tasks was that they were programmed without writing any custom perception or manipulation code, or doing any other technical work like running scripts. Each action was programmed and executed in five new scenes, different from the demonstration scene. Performing the tasks, including providing the demonstration and defining custom landmarks, took less than 30 minutes to do, each. The scenes, landmarks, programmed actions, and the test scenarios for each task are explained in Tables V-X.

Overall, our evaluation with the six tasks and 30 test scenarios showed that custom landmarks successfully enabled PbD actions to be anchored on arbitrary object or scene parts. The robot succeeded in 26 out of the 30 total scenes it was evaluated in (Tables V-X). Our search algorithm detected landmarks in cluttered or unconventional scenes, enabling PbD to be used in these arguably realistic scenarios.

TABLE VIII: Demo Task 4: Trash bin / recycling bin

Scene and objects: Trash and recycling bins in front of the robot; both bins have same dimensions, but the trash bin has a square opening, while the recycling bin has a circular opening. **Landmarks:** Square and circular openings of the bins **Actions:** Drop an item into either bin	
Test scenarios	***Success?***
Drop in trash bin, with both bins moved slightly	*Yes*
Drop in trash bin, with two bins swapped (Fig. 3(d))	*Yes*
Drop in recycling bin, with two bins rotated 45^o	*Yes*
Drop in trash bin, with only recycling bin in the scene	*Yes*
Drop in recycling bin, only recycling bin in the scene	*Yes*

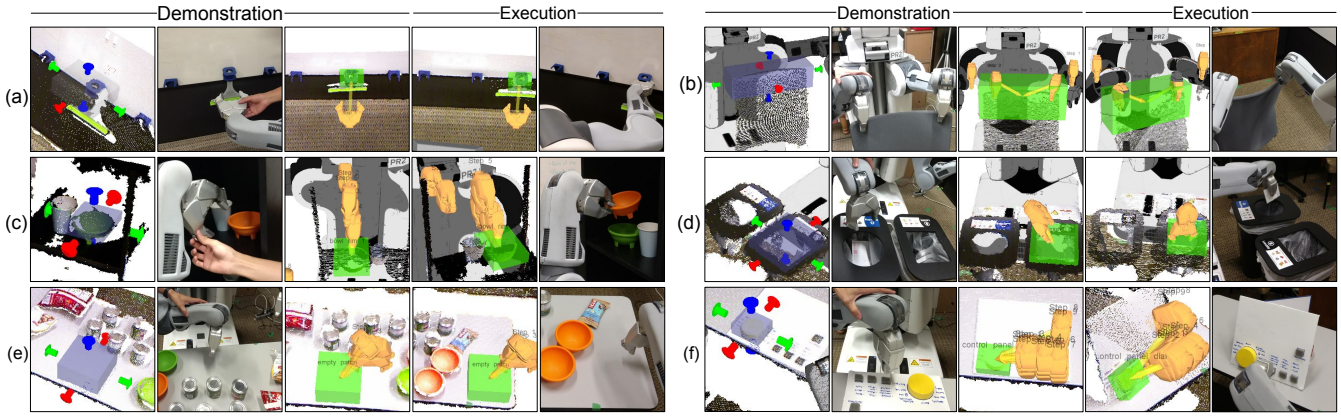

Fig. 3: PbD demonstrations with custom landmarks, as described in Section V. Each row (a)-(f) illustrates Tasks 1-6, respectively. The pictures in each row show, from left to right, 1) the box the user defines around the custom landmark of interest, 2) the user demonstrating an action relative to that landmark, 3) a visualization of the demonstrated action (orange grippers), and the custom landmark (green box), 4) a visualization of the PbD execution in a different scene, and 5) the robot executing the action in the new scene. (a) Shows the robot picking a Gripple tool, which was moved to a different position in one test. (b) Shows the robot grasping a slightly rotated chair. (c) Shows the robot picking the top bowl from a stack of bowls on a shelf. (d) Shows the robot placing an item in either a trash bin or a recycling bin. (e) Shows the robot finding an empty patch on a crowded table. (f) Shows the robot operating a mockup of a laundry control panel. During the demonstration, the control panel was flat on the table, but during one test, the control panel was upright.

TABLE IX: Demonstration Task 5: Crowded surface

Scene and objects: The robot searches for an empty spot on a crowded tabletop with various objects **Landmarks:** A flat tabletop patch **Actions:** Point to an empty spot	
Test scenarios	***Success?***
Different scenes with objects configured to vary the position of the empty patch. Two of the scenes had a box larger than the empty patch on the table.	*4/5*

TABLE X: Task 6: Control panel

Scene and objects: Control panel, similar to a laundry machine, with push buttons and a large central dial **Landmarks:** The central dial **Actions:** Pressing a particular sequence of buttons on the panel	
Test scenarios	***Success?***
Different scenes with the control panel position and incline orientation changed.	*4/5*

Our algorithm was robust enough to avoid actions from being executed when they were not supposed to. The robot refused to act in all cases where the necessary landmark was missing in the scene (Task 1: picking from empty rack, placing on full rack; Task 3: no bowls to pick up in scene with only cups, Task 4: no trash bin to drop in).

The failure cases were due to a few different reasons. In Task 1, we considered one of the correct executions to be a partial success, because a Gripple mount had to be manually supported to prevent it from being knocked down. This could have been avoided if the mount (fastened with Velcro) was attached more securely to the wall. In Task 5, our system thought that the surface of a flat box was an empty spot on the table. This shows, again, how our system depends on the uniqueness of the custom landmark's shape to work well. However, this could have been fixed if the robot compared the heights of the multiple empty patches it found. In Task 2, the angle of the chair was directly lined up with the robot's Kinect sensor, causing the appearance of the chair to change too much to localize. Similarly, in Task 6, the control panel was inclined at such an angle that the Kinect was unable to get full depth data for the dial to be localized. These last two failure cases show that using a single sensor and a single view of the scene can make our system brittle.

These tasks illustrate the advantageous properties of custom landmarks described in Section III-B:

1) Hanging tools (Task 1), grasping chairs (Task 2), picking bowls from a shelf (Task 3), and distinguishing trash bins (Task 4) all did not involve tabletop scenes.
2) Our system worked in the hard to segment shelf scene (Task 3), where bowls and cups were placed in contact.
3) Picking the top bowl from a stack of bowls (Task 3) showed that custom landmarks can be used to find partly occluded objects.
4) Picking and placing tools with Griples (Task 1) showed that custom landmarks can generalize across objects with the same part.
5) Most of the tasks would have required task-specific perception code without custom landmarks, even if an advanced object recognition system had been available.
6) Finally, the landmarks in most of the tasks did not represent objects. Instead they represented partial objects, parts of the scene, or even the lack of objects.

VI. Usability test

We also conducted a small-scale usability test of the bounding box interface, used to capture custom landmarks. 5 users were taught how to use our PbD system with custom landmarks, and were asked to program one of three tasks: putting away bowls on a table, retrieving one of three snacks, or stacking cylinders on top of each other. Participants created custom landmarks by moving the 6 sides of a 3D bounding box, as depicted in Fig. 3. After completing the task, users were asked in a survey to describe the most difficult part of specifying custom landmarks. We also recorded the time spent using the interface. Users were recruited from email lists associated with the University of Washington. Participants were required to not be robotics researchers and to have two years of programming experience. All participants were male and had an average age of 25.2 years.

Participants took an average of 1 minute 11 seconds (SD=42 seconds) to specify a custom landmark. All of the participants were able to use the bounding box interface to create viable landmarks for the corresponding PbD action. After completing the usability test, 2 of the 5 participants wrote that they had difficulty with the interface. One of these participants wrote that they wanted to have separate controls to translate the box, instead of adjusting the six sides of the box individually. The other of these participants wrote that they would like to have an automatic "guess" of the box, which could then be adjusted. The usability test suggests that, while the bounding box interface could be improved, it is functional enough for users to work with.

VII. Conclusion

This paper presents a way to enable end-users to customize the perception system of their robots, so as to support actions they would like to perform. We proposed representing task-relevant objects, object parts, or scene parts (*i.e.*, landmarks) in the robot's environment with point cloud patches that are specified by the user prior to programming an action. The user can then program manipulation actions by demonstrating end-effector poses relative to these landmarks. We presented an algorithm for localizing landmarks in new scenes. We systematically evaluated our algorithm with scene-landmark pairs and characterized its performance. Then, we showed real examples of complex manipulation tasks being programmed on a PR2 robot, and demonstrated that they work well in cluttered and unconventional (non-tabletop) scenes. Finally, we performed a usability test demonstrating the feasibility of our system for novice users.

References

[1] S. Alexandrova, M. Cakmak, K. Hsiao, and L. Takayama, "Robot programming by demonstration with interactive action visualizations," in *Robotics: Science and Systems (RSS)*, 2014, pp. 48–56.

[2] S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 5239–5246.

[3] J. Schulman, J. Ho, C. Lee, and P. Abbeel, "Learning from demonstrations through the use of non-rigid registration," *International Journal of Robotics Research*, 2013.

[4] S. Niekum, S. Chitta, A. G. Barto, B. Marthi, and S. Osentoski, "Incremental semantically grounded learning from demonstration." in *Robotics: Science and Systems*, 2013.

[5] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer Handbook of Robotics*. Springer, 2008, pp. 1371–1394.

[6] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[7] S. Chernova and A. L. Thomaz, "Robot learning from human teachers," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 8, no. 3, pp. 1–121, 2014.

[8] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *ICML*, vol. 97, 1997, pp. 12–20.

[9] S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert, "Learning movement primitives," in *International Symposium on Robotics Research.*, 2005, pp. 561–572.

[10] S. Calinon and A. Billard, "Statistical learning by imitation of competing constraints in joint space and task space," *Advanced Robotics*, vol. 23, no. 15, pp. 2059–2076, 2009.

[11] B. Akgun, M. Cakmak, J. W. Yoo, and A. L. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, 2012, pp. 391–398.

[12] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, 2009, pp. 763–768.

[13] M. Pardowitz, S. Knoop, R. Dillmann, and R. Zollner, "Incremental learning of tasks from user demonstrations, past experiences, and vocal comments," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 37, no. 2, pp. 322–332, 2007.

[14] S. Ekvall and D. Kragic, "Robot learning from demonstration: a task-level planning approach," *International Journal of Advanced Robotic Systems*, vol. 5, no. 3, pp. 223–234, 2008.

[15] H. B. Suay, R. Toris, and S. Chernova, "A practical comparison of three robot learning from demonstration algorithm," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 319–330, 2012.

[16] M. Wächter, S. Schulz, T. Asfour, E. Aksoy, F. Wörgötter, and R. Dillmann, "Action sequence reproduction based on automatic segmentation and object-action complexes," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2013, pp. 189–195.

[17] B. Akgun and A. Thomaz, "Simultaneously learning actions and goals from demonstration," *Autonomous Robots*, vol. 40, no. 2, pp. 211–227, 2016.

[18] M. Ehrennmann, D. Ambela, P. Steinaus, and R. Dillmann, "A comparison of four fast vision based object recognition methods for programming by demonstration applications," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 2, 2000, pp. 1862–1867.

[19] P. J. Besl and N. D. McKay, "Method for registration of 3-D shapes," in *Robotics-DL tentative*, 1992, pp. 586–606.

[20] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (FPFH) for 3D registration," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, 2009, pp. 3212–3217.

[21] A. G. Buch, D. Kraft, J.-K. Kamarainen, H. G. Petersen, and N. Krüger, "Pose estimation using local structure-specific shape and appearance context," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 2013, pp. 2080–2087.

[22] G. Anthes, "Deep learning comes of age," *Communications of the ACM*, vol. 56, no. 6, pp. 13–15, 2013.

[23] D. Maturana and S. Scherer, "VoxNet: A 3D convolutional neural network for real-time object recognition," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, 2015, pp. 922–928.

[24] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao, "3D ShapeNets: A deep representation for volumetric shapes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1912–1920.

[25] Z. Xu and M. Cakmak, "Enhanced robotic cleaning with a low-cost tool attachment," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, 2014, pp. 2595–2601.